

FREE ANNOTATED DATA FOR DEEP LEARNING IN MICROSCOPY? A HITCHHIKER'S GUIDE

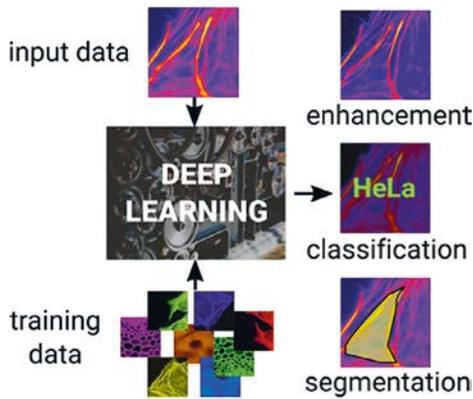
Adrian SHAJKOFCI^{1,2}, Michael LIEBLING^{1,3*}

¹ Idiap Research Institute, Rue Marconi 19, 1920 Martigny, Switzerland

² École Polytechnique Fédérale de Lausanne (EPFL), Route Cantonale, 1015 Lausanne, Switzerland

³ University of California, Santa Barbara, CA 93106-9010, USA

*michael.liebling@idiap.ch



In microscopy, the time burden and cost of acquiring and annotating large datasets that many deep learning models take as a prerequisite, often appears to make these methods impractical. Can this requirement for annotated data be relaxed? Is it possible to borrow the knowledge gathered from datasets in other application fields and leverage it for microscopy? Here, we aim to provide an overview of methods that have recently emerged to successfully train learning-based methods in bio-microscopy.

<https://doi.org/10.1051/photon/202010430>

Optical microscopy, despite being an invaluable tool in biology and medicine to observe and quantify cellular function, organ development, or disease mechanisms, requires constant trade-offs between spatial, temporal, and spectral resolution, invasiveness, acquisition time, and post-processing effort. As for other imaging fields, learning-based techniques are having a major impact in microscopy, where their potential to improve resolution, reduce invasiveness, or increase the speed of microscopy acquisitions has recently been demonstrated. Yet despite the ever-increasing computational power, it is often the lack of labeled training data that is the limiting factor for wide adoption. Annotating data is

often a lengthy and expensive task, since it involves tedious work, generally by skilled experts. Annotation can be especially challenging in the case of three-dimensional images, common in microscopy, despite the development tools dedicated to this task. While the acquisition and annotation of volumetric data is common in some medical settings and for certain modalities because of their wide use for healthcare applications (e.g. magnetic resonance imaging (MRI), annotated by trained radiologists), microscopy applications often lack similarly large and high-quality annotated volumetric datasets.

For many applications in biology, the time burden, cost, or physical feasibility of acquiring and annotating datasets for deep learning models de novo is simply out of the question. Can this requirement for annotated

data be relaxed? Is it possible to borrow the knowledge gathered from datasets in other application fields, such as e-commerce or computer gaming, and leverage it for bio-microscopy? Specifically, could annotated datasets be generated from realistic synthetic models of tridimensional objects? Or could more abstract prior knowledge about the data be used to enhance the resolution? Here, we aim to provide an overview of some solutions that have emerged to tackle the problem of gathering sufficient annotated data to train learning-based methods in bio-microscopy. We have grouped the approaches in four broad categories: developing manual annotation strategies, learning from annotated images from other domains, building annotations from simulations, and using self-annotated data. This quest for annotated data is summarized in Figure 1.

OVERVIEW OF DATA ANNOTATION STRATEGIES FOR DEEP LEARNING METHODS IN BIOMICROSCOPY

LEARNING FROM LARGE TRAINING DATASETS

Image segmentation is a common computer vision task in which each pixel of an image is assigned a label corresponding to the object it belongs to. In microscopy, it is used, for example, to delineate, identify, and count cells. Learning-based segmentation methods require a training data set composed of images together with masks that correspond to the objects and their coordinates in the images. This annotation is usually performed manually and can be tedious. Other fields have integrated the task of annotating images into security forms on the web (to exclude robots from accessing content) or into entertaining game puzzles, such as to entice the public to provide quality annotations. Sullivan *et al.* have proposed to annotate microscopy data by a similar approach, as part of a multiplayer computer game named Eve Online [1]. Using the publicly available data set from the Cell Atlas of the Human Protein Atlas (HPA), they obtained, over one year, nearly 33 million classifications of subcellular localization patterns of immunostained proteins in 20 different organelles and cellular structures. The results were successfully used as a training dataset for a segmentation deep neural network (DNN).

LEARNING FROM OTHER DOMAINS

Training a machine learning algorithm by re-using computer vision data sets that were originally intended for other tasks is at the core of transfer learning. This approach can overcome the lack of annotated data in one field (such as microscopy) if annotated data exist in a different field. A wide variety of annotated image datasets are available, such as ImageNet, MS-COCO or Places, which contain foremost scenes depicting everyday objects and situations. In addition to allowing access to a large number of examples, learning from natural-scene images can help avoid learning on images that contain unwanted aberrations, such as blur and low-light noise commonly found in microscopy images.

It becomes essential when developing methods that aim at removing aberrations, for example when using a DNN pipeline for deconvolution or denoising of microscopy images, since the ground truth of natural-scene photographs is more readily accessible than fine structures that challenge the resolution limits of the most powerful microscopes.

In most cases of transfer learning, pre-trained models are used as feature extractors, as the input of an unsupervised classifier such as support vector machines (SVM), or are fine-tuned with job-specific training data. However, there are some applications where transfer learning can be used without fine-tuning, specifically, when the feature space can be mapped identically in both domains. For image deconvolution and depth estimation, a description of the Point Spread Function (PSF) is a pre-requisite and accurately determining PSF parameters is therefore essential. Recently, we formulated the problem of retrieving the physical PSF parameters of an optical microscope as a regression task for which we trained a DNN. Interestingly, when learning from data that consist of textured images, even if they are different from the end-use application, the trained model remains just as accurate as when it is trained on microscopy images [2]. The ability to train on generic data also helps prevent over-fitting the trained model to a specialized and narrow data set, which would make the model less suitable to be used in other situations. It also suggests the possibility of generalizing the trained network to other data types, provided that the new data and the data used for training share a common feature space.

LEARNING FROM SIMULATED DATA SETS

Generating annotated data from simulations is another effective approach to produce large and reliable training data sets. In some computer vision applications, such as autonomous car driving, data from simulated computer graphics 3D environments have been effectively used to train segmentation methods. Examples of such datasets include the Flying chairs dataset, or the GTA5 dataset derived from a computer game.



Automatisez votre système d'alignement optique



MPM-211: 4 ports optiques
MPM-212: 2 ports optiques
MPM-213: 4 ports courant



MPM-210: Le chassis

avec le wattmètre optique MPM et sa mesure de courant intégrée

- Intégré avec nos sources lasers
- Jusqu'à 20 canaux
- Plage dynamique > 80 dB



Visionnez la vidéo ici

www.santec.com

santec-emea@santec.com

Generally, the accuracy of DNNs trained only on simulated data is poor, due to the extreme complexity of natural scenes that can hardly be reproduced with simple simulations.

In microscopy, image complexity remains fairly low in some modalities such as single-molecule localization microscopy (SMLM), where features consist of dots. The image processing task, for which DNNs have been used, consists in converting images of random subsets of activated fluorophores, obtained over many consecutive diffraction-limited frames, into a high-precision point cloud. Data simulation of a realistic diffraction-limited ground truth is achievable, for example, by filtering the expected objects by the optical PSF and take into account a realistic noise model [3]. Such approaches have allowed recent DNN methods such as [4] to recover densely overlapping PSFs of multiple emitters over a large axial range and output a list of their 3D positions. The corresponding training data set were created by simulating a large number of images using randomly generated 3D patterns and the phase mask that governs the PSFs modelled on the physical implementation of the microscope. In order for the simulations to work as a sole training data (*i.e.* the output of the network) by closely modeling the data generation model to the physical properties of the optical system ([2], [4]).

LEARNING FROM THE INPUT IMAGE ITSELF

DNNs are also used for image enhancement such as deblurring (going from a blurry image to a sharper image), denoising (going from a noisy to a clean image), and for super-resolution (going from a low-resolution image to a higher-resolution image). These networks are traditionally trained on pairs of clean and distorted images. In microscopy, where the raw images often already reach the physical limits due to diffraction, a higher-resolution ground-truth image

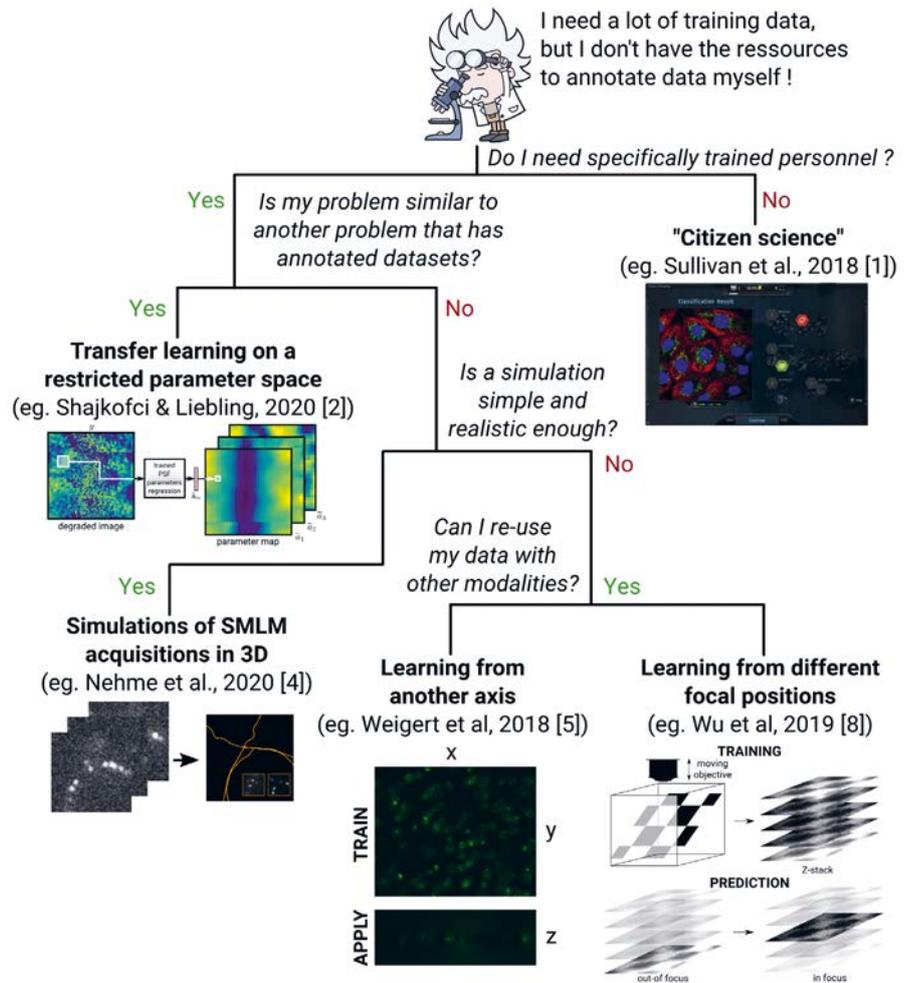


Figure 1. The quest for less painful annotation: overview of questions and methods. The illustrations were made by AS and adapted from the original papers.

is simply not accessible. To cope with this problem, it is sometimes possible to use approaches that exploit particular features of the data, such as its isotropy or the availability of complementary imaging modalities.

Weigert *et al.* [5] proposed a pipeline aimed at restoring images using semi-synthetic training data. Specifically, it restores the axial resolution of volumetric images lost due to the axial elongation of the optical PSF and the low axial sampling rate. By assuming that similar features are to be expected regardless of sample orientation, the method leverages the fact that these features can be much better resolved in lateral views than in depth, hence the training to improve depth-resolution is done based on

I need a lot of training data, but I don't have the resources to annotate data myself!



Do I need specifically trained personnel?

the latter images. Nevertheless, for many other applications, access to higher quality images or synthetic data is not available. To overcome this limitation, Krull *et al.* [6] developed a self-supervised training method for denoising based on the assumption that the noise is independent from pixel to pixel and that the true intensity of a pixel can be predicted from the local image context since it is not locally independent. The method involves a noise model whose probability distribution is learned from the training data and a network is trained to discriminate the underlying image from the noise. As a training set, a very small dataset of noisy images of the same type can be used. This method is virtually equivalent to training a DNN for every specific noise distribution.

A similar training scheme was also used in [7, 8], where the authors trained a generative adversarial network (GAN)-based DNN to transform an acquired low-resolution image into a

high-resolution image using matched pairs of experimentally acquired images after registration and alignment. In [7], these pairs came from images of the same object using a confocal microscope and a STED microscope. In [8], the authors used a similar DNN to generate images that look as if they had been taken from another focal plane by training from images acquired at different heights in the sample. For both applications, the authors caution that the network must be (re)-trained for each specific image modality or experimental setup, as the methods do not produce ideal results otherwise. The application of such methods therefore remains somewhat limited to cases where the type of images and microscope settings are known beforehand and where a high number of similar images can be acquired, which could be particularly relevant for time-lapse imaging, high-resolution 3D stacks, or imaging of histological samples prepared under controlled and standardized conditions.

CONCLUSIONS

Deep learning technologies have enabled multiple applications that are transforming our day-to-day routines, including the way we approach

microscopy. While limitations such as network capacity (can networks learn to predict from the wide variety of data types common in microscopy?), generalization (can a network trained on one type of data be used to handle other types of data?), and overfitting (is the network limited to predicting only what it has already seen?) are some pressing issues that the field is facing, the lack of good quality training data is likely the single most important aspect that affects accuracy and effectiveness of tasks such as enhancement, classification or segmentation. The most promising methods to overcome the scarcity of training data appear to leverage prior

knowledge of the physical objects and image formation process [4], [5], [2], or of the noise distribution [6].

Even if we are still a long way from a blind pipeline that will enhance, classify or segment biological data without tedious annotation work, good knowledge of the problem and assumptions about the data allow scientist to already reap the benefits of deep learning tools by crafting adapted training sets without having to produce or wait for the availability of large annotated sets. ●

ACKNOWLEDGEMENT

This work was supported by the Swiss National Science Foundation, Grant 206021.

REFERENCES

- [1] D.P. Sullivan *et al.*, *Nat. Biotechnol.* **36**, 820 (2018)
- [2] A. Shajkofci, M. Liebling, *IEEE Trans. Image Process.* **29**, 5848 (2020)
- [3] D. Sage *et al.*, *Nat. Methods* **16**, 387 (2019)
- [4] E. Nehme *et al.*, *Nat. Methods* **17**, 734 (2020)
- [5] M. Weigert *et al.*, *Nat. Methods* **15**, 1090 (2018)
- [6] A. Krull *et al.*, *Front. Comput. Sci.* **2**, 5 (2020)
- [7] H. Wang *et al.*, *Nat. Methods* **16**, 103 (2019)
- [8] Y. Wu *et al.*, *Nat. Methods* **16**, 1323 (2019)

SPECTROGON

State of the art products

Interference filters

- 200 to 15000 nm
- Bandpass
 - Longwave-pass
 - Shortwave-pass
 - Broad-bandpass
 - Neutral density
 - Web stock items



Holographic gratings

- 150 to 2000 nm
- Pulse compression
 - Telecom
 - Laser tuning
 - Monochromator
 - Spectroscopy
 - Web stock items



UK: sales.uk@spectrogon.com • Tel +44 1592770000
 Sweden (headquarters): sales.se@spectrogon.com • Tel +46 86382800
 US: sales.us@spectrogon.com • Tel +1 9733311191

www.spectrogon.com